

Social Teaching: Being Informative vs. Being Right in Sequential Decision Making

Joong Bum Rhim and Vivek K Goyal

Abstract—We show that it can be suboptimal for Bayesian decision-making agents employing social learning to use correct prior probabilities as their initial beliefs. We consider sequential Bayesian binary hypothesis testing where each individual agent makes a binary decision based on an initial belief, a private signal, and the decisions of all earlier-acting agents—with the actions of precedent agents causing updates of the initial belief. Each agent acts to minimize Bayes risk, with all agents sharing the same Bayes costs for Type I (false alarm) and Type II (missed detection) errors. The effect of the set of initial beliefs on the decision-making performance of the last agent is studied. The last agent makes the best decision when the initial beliefs are inaccurate. When the private signals are described by Gaussian likelihoods, the optimal initial beliefs are not haphazard but rather follow a systematic pattern: the earlier-acting agents should act as if the prior probability is larger than it is in reality when the true prior probability is small, and vice versa. We interpret this as being *open minded* toward the unlikely hypothesis. The early-acting agents face a trade-off between making a correct decision and being maximally informative to the later-acting agents.

Index Terms—Bayesian hypothesis testing, distributed detection, human decision making, likelihood ratio tests, sequential decision making, social learning, social networks, team theory.

I. INTRODUCTION

Consider decision-making agents facing the task of choosing between two alternatives. Each agent has a private signal, which is not visible to the other agents. The agents sequentially make their individual decisions, which are visible to other agents. An agent's action contains some information about the right (or better) choice, so subsequent agents can learn from the action and reflect it in their own actions. For example, when you want to choose between two alternatives when buying a new phone, the choices made by your colleagues can affect your judgment.

Being influenced by earlier-acting agents has been termed *social learning* [1]. It has generally been studied in settings where each agent has no motivation beyond making a correct choice for himself. In this paper, we study the effect of an agent's action on subsequent agents and find that making correct decisions is generally not equivalent to providing information to other agents that maximally benefits them in their decision making tasks. Accounting for the effect on other agents could be termed *social teaching*. Also, in any scenario with social learning, the earlier-acting agents can be seen as *advisers* to the later-acting agents. As will be detailed later, we find that a good adviser should be *open minded* in the sense

of being more receptive to the *a priori* less likely alternative than she would have been if she were only interested in being right rather than also interested in being informative.

The framework of sequential decision making with social learning was independently introduced in [2] and [3]. These works focused primarily on *herding*, which is for all agents beyond some index to take the same action. They showed that an incorrect herd would arise with positive probability when private signals are boundedly informative.¹ For example, the private signals were assumed to be binary and to give true or false information, each with positive probability. It can happen that a couple of the first agents receive false private signals and thus choose wrong actions. Then the effect of these actions on the beliefs of subsequent agents can be so great as to cause them to ignore their private signals and follow their precedent agents. The private signals are bounded so that they cannot be strong enough to overcome the effect of the wrong actions.

Subsequently, [4] showed that learning is incomplete—beliefs are not eventually focused on the true state—if private signals are boundedly informative, but agents will asymptotically settle on the optimal action otherwise. Recently, [5] extends the result to general network topologies where each agent can observe decisions made by its neighbors instead of all previous agents.

In another related line of work, [6] studies the effect of social learning in a quickest detection problem, in which agents keep updating their beliefs based on previous decisions and detect the time at which an underlying state changes. It has a similar framework to [7], which studied update of private information in a finite memory.

This paper differentiates itself from the literature in that it considers unbounded private signals and does not focus on herding behavior. In addition, we focus largely on the effect of prior probabilities in decision making. We do not assume that an agent knows a correct prior probability for the decision at hand. Even if he does, we do not assume he takes the shortsighted approach of using the prior probability only to optimize the correctness of his own decision. Instead, we study the effect of the prior probability on the decisions of subsequent agents. Sequential decision making is considered from a signal-processing perspective in [8] as well. Its model is similar to ours except that there all agents know the true prior probability; this difference changes the problem substantially.

Our criterion for optimality is the Bayes risk of the final agent; we assume a sequential decision making model in

This material is based upon work supported by the National Science Foundation under Grant No. 1101147.

The authors are with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

¹A signal Y generated under a state H is called *boundedly informative* if there exists $\kappa > 0$ such that $\kappa < f_{Y|H}(y|h) < 1/\kappa$ for all y and h .

which only the decision made by the final agent matters. Since sequential decision making is a hypothesis testing problem, agents adopt likelihood ratio tests to choose their actions [9]. As they observe decisions or actions chosen by precedent agents, they compute or *update* their beliefs to perform more precise likelihood ratio tests. The update process depends on the initial prior belief and the history of decisions. We derive a recursive belief update function.

Bayes-optimal agents need to know the prior probability in order to perform the likelihood ratio test. Hence, intuitively, agents with wrong prior beliefs should degrade the decision making and yield higher Bayes risk. In addition, they would misunderstand the public signals because they do not know others' beliefs.

Contrary to intuition, it turns out that wrong beliefs may improve the decision made by the final agent in sequential decision making. Especially when the private signals are distorted by additive Gaussian noise, the optimal first agent is open minded: He acts as if the prior probability is larger than it is in reality when the true prior probability is small, and vice versa.

Section II provides additional background and motivation from human decision makers. Section III describes our sequential decision making model. In Section IV, we investigate how agents interpret the decision history, update their beliefs, and make decisions according to their positions in the chain of agents. It is proven in Section V that the true prior probability is not the optimal prior belief for $N = 2$. Examples for Gaussian likelihoods are presented in Section VI. Section VII concludes the paper.

II. BACKGROUND

The mathematical model presented herein abstracts human decision makers so as to be broadly applicable. We are motivated in part by a study about the correlation between a defendant's physical appearance and juror decisions [10]. It states that jurors feel a defendant more intelligent when the defendant is wearing eyeglasses, which leads to fewer guilty verdicts. It also says that wearing eyeglasses is especially effective for African-American defendants. Several other studies have also revealed the importance of a defendant's physical appearance on a jury's decision making [11]–[13]. From a logical standpoint, a defendant's appearance should be irrelevant to judicial decisions because eyeglasses have nothing to do with crimes, and we believe that jurors do their best to make fair and reasonable decisions based only on evidence. Then why does this happen in reality?

We find the answer in Bayesian reasoning and the concept of prior probability. Let us liken a jury trial to a hypothesis testing problem. A defendant is metaphorically an object in one of two states: guilty or not guilty. Perceptions of evidence presented by a prosecutor or defense counsel are noisy observations about the defendant's true state. Jurors are detectors that make a decision based upon the noisy observations. However, one element of hypothesis testing is missing: the prior probability that the defendant commits a crime.

Reasonable human beings resemble Bayesian decision makers [14]–[17]; they need to know the prior to compute the




Fig. 1. A sequential decision making model with N agents (Alexis, Blake, ..., Norah). The n th agent can observe $n-1$ decisions made by the precedent agents.

posterior probability. The prior probability can be critical to the verdict when the evidence at trial is ambiguous. The problem is that the jurors cannot know the defendant's true prior probability. Hence, before reaching a verdict, the jurors judge the defendant's prior probability by how intelligent, how attractive, how friendly, and how threatening the defendant "looks." They may be able to estimate the prior probability close to the true value but their estimates would not be the same as the defendant's true prior probability.

We are not defending or criticizing this phenomenon but just focusing on an interesting issue raised by it: Human agents perform Bayesian hypothesis testing with inaccurate knowledge of prior probabilities. While the prior probability is one of the basic elements of estimation, the effect of accuracy of the prior probability has not received a great deal of attention. Initially building upon [18], we have previously studied the effect of categorization of problem instances as inducing quantization of prior probabilities [19]–[22]. The present paper is more fundamental in that it addresses whether accurate prior probabilities are even the most favorable. We have found that inaccurate perception of the prior probability may be beneficial in sequential decision making.

III. PROBLEM DESCRIPTION

Consider the sequential decision making model depicted in Fig. 1. There is an object in a binary state $H \in \{0, 1\}$ with probability $\mathbb{P}(\{H = 0\}) = p_0$ and $\mathbb{P}(\{H = 1\}) = 1 - p_0$. There are also N agents that sequentially detect the state. Agents do not know the true prior probability of the object. Instead, the n th agent perceives it as q_n . The n th agent observes decisions made by precedent agents, $\{\hat{H}_1, \dots, \hat{H}_{n-1}\}$, as well as a signal about H , Y_n , generated from a likelihood function $f_{Y_n|H}$.² The private signals $\{Y_n\}$ are conditionally independent given H and are identically distributed. We assume that the likelihood ratio $f_{Y_n|H}(y_n|1)/f_{Y_n|H}(y_n|0)$ is an increasing function of y_n .

The n th agent can extract some information from the $n-1$ precedent decisions. The decisions are, however, biased by the false impressions that the precedent agents have of the object. Even worse, the n th agent does not know what q_1, \dots, q_{n-1} are. Thus, the n th agent assumes that they are all equal to q_n and interprets the precedent decisions accordingly. Using the history of decisions, the n th agent updates its prior belief

²The signal in this model has a continuous value while the corrupted signal has a discrete value in the initially-proposed framework [2], [3].

before applying its likelihood ratio test. We will define a recursive function that describes the belief update in Section IV.³

Our interest is in the last agent Norah and her decision \hat{H}_N . Upon observing her private signal Y_N and the $N-1$ precedent decisions $\hat{H}_1, \dots, \hat{H}_{N-1}$, she determines her decision rule. We evaluate the decision rule by a common criterion, Bayes risk, which measures the expected cost of her decision. The relative importance of correct decisions and errors can be abstracted as a cost function $c(\hat{H}, H)$, which defines penalties for false alarm or Type I error (choosing $\hat{H} = 1$ when $H = 0$), correct rejection (choosing $\hat{H} = 0$ when $H = 0$), hit (choosing $\hat{H} = 1$ when $H = 1$), and missed detection or Type II error (choosing $\hat{H} = 0$ when $H = 1$). For simplicity, we assume the correct decisions have zero cost and use the shorthand notations $c_{10} = c(1, 0)$ and $c_{01} = c(0, 1)$ respectively for costs of false alarms and missed detections. In addition, we consider agents have the same costs; they are a team in the sense of Radner [23]. Then the Bayes risk is given by

$$R_N = c_{10}p_0p_{\hat{H}_N|H}(1|0) + c_{01}(1-p_0)p_{\hat{H}_N|H}(0|1). \quad (1)$$

The computation of (1) depends on the previous decisions $\hat{H}_1, \dots, \hat{H}_{N-1}$. Therefore, the correct computation of the expected cost is

$$R_N = \sum_{\hat{h}_1, \dots, \hat{h}_{N-1}} \left(c_{10}p_0p_{\hat{H}_N, \hat{H}_{N-1}, \dots, \hat{H}_1|H}(1, \hat{h}_{N-1}, \dots, \hat{h}_1|0) + c_{01}(1-p_0)p_{\hat{H}_N, \hat{H}_{N-1}, \dots, \hat{H}_1|H}(0, \hat{h}_{N-1}, \dots, \hat{h}_1|1) \right). \quad (2)$$

We will discuss the optimal values of q_n that minimize (2).

It is important to note that, in our model, each agent uses a decision rule optimized for her own belief; the agents do not adjust their decision rules for the sake of Norah. In other words, for all $n = 1, \dots, N$, the n th agent adopts the decision rule that minimizes her Bayes risk R_n , and her decision is shown to the other agents as a public signal. In contrast, the agents could adjust their decision rules in an attempt to minimize the Bayes risk of a single collective decision, as studied for the combination of social learning and aggregation by voting in [24].

We now introduce additional notation for the rest of the paper. Random variables are in uppercase while their realizations are in lowercase. We denote a probability density function (pdf) of a continuous random variable as f and a probability mass function (pmf) of a discrete random variable as p . A subscript number n means “of the n th agent.” Superscript alphabet A (B) means “upon observing Alexis’s (Blake’s) decision”; we sometimes use 0 or 1 instead of the Roman alphabet to specify a decision value. For example, q_3^{AB} denotes the updated belief of the third agent, Chuck, upon observing Alexis’s and Blake’s decisions \hat{H}_1 and \hat{H}_2 , and q_3^{10} denotes Chuck’s updated belief upon observing $\hat{H}_1 = 1$ and $\hat{H}_2 = 0$. Subscript alphabet A (B) means “that Alexis (Blake) thinks.”

³We use the term *belief* to distinguish from the true prior probability and to capture that it is what agents believe as the prior probability. Agents initially perceive the prior probability in some way, which we call the *prior beliefs*. After they observe precedent decisions, they modify the prior beliefs. The beliefs are then called *updated*.

For example, Blake thinks that the probability of Alexis choosing 0 when the true state is 0 is $p_{\hat{H}_1|H}(0|0)_B$. We need to clarify who thinks it because the agents are not aware of others’ prior beliefs. This will be explained in detail in Section IV-B.

IV. PRIOR BELIEF UPDATE AND DECISION MAKING

Our model assumes unbounded private signals. Thus, unlike in [2], [3], it is always possible that a subsequent agent may not follow previous decisions; that is, incorrect herding does not occur. We now discuss the utilization of a decision history as well as private signals for Bayesian hypothesis testing. A position-wise decision-making strategy will be provided. This can be interpreted as each agent updating his prior belief based on the decision history and then applying a likelihood ratio test with his private signal.

A. Alexis, the First Agent

Alexis performs normal binary hypothesis testing because she has no precedent decision. She use the following likelihood ratio test with her prior belief q_1 :

$$\frac{f_{Y_1|H}(y_1|1)}{f_{Y_1|H}(y_1|0)} \underset{\hat{H}_1(y_1)=0}{\overset{\hat{H}_1(y_1)=1}{\gtrless}} \frac{c_{10}q_1}{c_{01}(1-q_1)}. \quad (3)$$

Since the likelihood ratio is increasing in y_1 , the likelihood ratio test can be simplified to comparison with an appropriate decision threshold:

$$y_1 \underset{\hat{H}_1(y_1)=0}{\overset{\hat{H}_1(y_1)=1}{\gtrless}} \lambda(q_1), \quad (4)$$

where $\lambda(q)$ denotes the decision threshold that satisfies

$$\frac{f_{Y|H}(\lambda|1)}{f_{Y|H}(\lambda|0)} = \frac{c_{10}q}{c_{01}(1-q)}. \quad (5)$$

B. Blake, the Second Agent

Blake observes Alexis’s decision. Considering \hat{H}_1 as another corrupted signal of H like Y_2 , he modifies the likelihood function (3) with his prior belief q_2 to

$$\frac{f_{Y_2, \hat{H}_1|H}(y_2, \hat{h}_1|1)}{f_{Y_2, \hat{H}_1|H}(y_2, \hat{h}_1|0)} \underset{\hat{H}_2(y_2)=0}{\overset{\hat{H}_2(y_2)=1}{\gtrless}} \frac{c_{10}q_2}{c_{01}(1-q_2)}. \quad (6)$$

In the left-hand side of (6),

$$f_{Y_2, \hat{H}_1|H}(y_2, \hat{h}_1|h) = f_{Y_2|H}(y_2|h)p_{\hat{H}_1|H}(\hat{h}_1|h)$$

because the private signals Y_1 and Y_2 are conditionally independent given H . We can rewrite (6) as⁴

$$\frac{f_{Y_2|H}(y_2|1)}{f_{Y_2|H}(y_2|0)} \underset{\hat{H}_2(y_2)=0}{\overset{\hat{H}_2(y_2)=1}{\gtrless}} \frac{c_{10}q_2}{c_{01}(1-q_2)} \frac{p_{\hat{H}_1|H}(\hat{h}_1|0)_B}{p_{\hat{H}_1|H}(\hat{h}_1|1)_B}. \quad (7)$$

⁴The subscript “B” in the term “ $p_{\hat{H}_1|H}(\hat{h}_1|h)_B$ ” indicates the value of $p_{\hat{H}_1|H}(\hat{h}_1|h)$ that Blake (the second agent) thinks. We specify this because Blake does not know Alexis’s belief q_1 . Thus, he interprets her decision based on his belief q_2 . The value is different from the true value of $p_{\hat{H}_1|H}(\hat{h}_1|h) = p_{\hat{H}_1|H}(\hat{h}_1|h)_A$. Of course, it will be also different from what Chuck—the third agent—thinks, which is denoted by $p_{\hat{H}_1|H}(\hat{h}_1|h)_C$. This will be explained in the next paragraph.

The likelihood ratio test (7) can be interpreted with Blake updating his prior belief upon observing Alexis's decision \widehat{H}_1 . Combined with q_2 , $p_{\widehat{H}_1|H}(\widehat{h}_1|h)_2$ updates his prior belief from q_2 to q_2^Δ :

$$\frac{q_2^\Delta}{1 - q_2^\Delta} = \frac{q_2}{1 - q_2} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_B}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_B}. \quad (8)$$

The updated belief is

$$\begin{aligned} q_2^\Delta &= \frac{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)_B}{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)_B + (1 - q_2) p_{\widehat{H}_1|H}(\widehat{h}_1|1)_B} \\ &= \frac{p_{\widehat{H}_1,H}(\widehat{h}_1,0)_B}{p_{\widehat{H}_1,H}(\widehat{h}_1,0)_B + p_{\widehat{H}_1,H}(\widehat{h}_1,1)_B} \\ &= p_{H|\widehat{H}_1}(0|\widehat{h}_1)_B. \end{aligned} \quad (9)$$

We have to make clear that Blake does not correctly compute $p_{\widehat{H}_1|H}(\widehat{h}_1|h)_B$ because he does not know q_1 . The true probability is given by

$$\begin{aligned} p_{\widehat{H}_1|H}(0|h) &= p_{\widehat{H}_1|H}(0|h)_A = \mathbb{P}(\{Y_1 \leq \lambda(q_1) | H = h\}) \\ &= \int_{-\infty}^{\lambda(q_1)} f_{Y|H}(y|h) dy, \end{aligned} \quad (10)$$

but Blake evaluates Alexis's decision \widehat{H}_1 as if it were made based on q_2 not q_1 :

$$\begin{aligned} p_{\widehat{H}_1|H}(0|h)_B &= \mathbb{P}(\{Y_1 \leq \lambda(q_2) | H = h\}) \\ &= \int_{-\infty}^{\lambda(q_2)} f_{Y|H}(y|h) dy. \end{aligned} \quad (11)$$

An interesting observation is that Alexis's biased belief q_1 does not affect Blake's belief update. There is no trace of q_1 in (9) and (11). Suppose that Alexis knows true prior probability p_0 and uses the decision threshold $\lambda(p_0)$. Still Blake, who does not know what belief Alexis has, thinks that the conditional probability of Alexis declaring $\widehat{H}_1 = 0$ is given by (11) and updates his belief as in (9). It is clear in (9) that the updated belief depends only on Blake's initial belief and Alexis's decision.

However, Alexis's prior belief still affects Blake's performance in some way, which is related to the probability of error. Alexis's biased belief changes the probability of her decision. The changed probability is embedded in the probability of Blake's decision:

$$\begin{aligned} p_{\widehat{H}_2|H}(\widehat{h}_2|0) &= \sum_{\widehat{h}_1} p_{\widehat{H}_2,\widehat{H}_1|H}(\widehat{h}_2,\widehat{h}_1|0) \\ &= p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|0,0)_B \times p_{\widehat{H}_1|H}(0|0)_A \\ &\quad + p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|1,0)_B \times p_{\widehat{H}_1|H}(1|0)_A, \\ p_{\widehat{H}_2|H}(\widehat{h}_2|1) &= \sum_{\widehat{h}_1} p_{\widehat{H}_2,\widehat{H}_1|H}(\widehat{h}_2,\widehat{h}_1|1) \\ &= p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|0,1)_B \times p_{\widehat{H}_1|H}(0|1)_A \\ &\quad + p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|1,1)_B \times p_{\widehat{H}_1|H}(1|1)_A. \end{aligned}$$

Thus, Alexis's biased belief changes the probability of Blake's decision as well as that of her decision.

C. Chuck, the Third Agent

Chuck's detection process is the same as Blake's. He observes both Alexis's and Blake's decisions and also updates his prior belief q_3 like in (8):

$$\begin{aligned} \frac{q_3^{\Delta B}}{1 - q_3^{\Delta B}} &= \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_2,\widehat{H}_1|H}(\widehat{h}_2,\widehat{h}_1|0)_C}{p_{\widehat{H}_2,\widehat{H}_1|H}(\widehat{h}_2,\widehat{h}_1|1)_C} \\ &= \left(\frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_C}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_C} \right) \frac{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,0)_C}{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,1)_C}. \end{aligned}$$

Be careful that \widehat{H}_1 and \widehat{H}_2 are not conditionally independent given H because Blake's decision \widehat{H}_2 depends on Alexis's decision \widehat{H}_1 :

$$p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,0) \neq p_{\widehat{H}_2|H}(\widehat{h}_2|0).$$

Chuck's update process can be split into in two steps. The first step is to infer Blake's updated belief based on Alexis's decision:

$$\frac{q_3^\Delta}{1 - q_3^\Delta} = \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_C}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_C}. \quad (12)$$

The second step is to update his own belief from q_3^Δ based on Blake's decision:

$$\frac{q_3^{\Delta B}}{1 - q_3^{\Delta B}} = \frac{q_3^\Delta}{1 - q_3^\Delta} \frac{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,0)_C}{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,1)_C}. \quad (13)$$

Please note that Chuck does not know Alexis's and Blake's prior beliefs, q_1 and q_2 , like Blake did not know Alexis's. Thus Chuck infers everything based on his own belief q_3 , which is indicated by the subscript "C" in (12) and (13).

Details of computations of (12) and (13) are as follows:

$$p_{\widehat{H}_1|H}(0|h)_C = \int_{-\infty}^{\lambda(q_3)} f_{Y_1|H}(y|h) dy, \quad (14a)$$

$$p_{\widehat{H}_1|H}(1|h)_C = \int_{\lambda(q_3)}^{\infty} f_{Y_1|H}(y|h) dy. \quad (14b)$$

Substituting (14) in (12), Chuck can compute q_3^Δ for $\widehat{H}_1 = 0$ and $\widehat{H}_1 = 1$ respectively:

$$q_3^0 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{-\infty}^{\lambda(q_3)} f_{Y_1|H}(y|1) dy}{\int_{-\infty}^{\lambda(q_3)} f_{Y_1|H}(y|0) dy}}, \quad (15a)$$

$$q_3^1 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{\lambda(q_3)}^{\infty} f_{Y_1|H}(y|1) dy}{\int_{\lambda(q_3)}^{\infty} f_{Y_1|H}(y|0) dy}}. \quad (15b)$$

Then,

$$\begin{aligned} p_{\widehat{H}_2|\widehat{H}_1,H}(0|\widehat{h}_1,h)_3 &= \mathbb{P}(\{Y_2 \leq \lambda(q_3^\Delta) | H = h\}) \\ &= \int_{-\infty}^{\lambda(q_3^\Delta)} f_{Y_2|H}(y|h) dy, \end{aligned} \quad (16a)$$

$$\begin{aligned} p_{\widehat{H}_2|\widehat{H}_1,H}(1|\widehat{h}_1,h)_3 &= \mathbb{P}(\{Y_2 > \lambda(q_3^\Delta) | H = h\}) \\ &= \int_{\lambda(q_3^\Delta)}^{\infty} f_{Y_2|H}(y|h) dy. \end{aligned} \quad (16b)$$

Even though the value of \widehat{h}_1 may not seem to be used in (16), it is inherent in q_3^Δ and affects the computation results. Chuck's updated belief $q_3^{\Delta B}$ is obtained by substituting (15) and (16) in (13).




Fig. 2. The function $U_4(q_4, \hat{h}_1, \hat{h}_2, \hat{h}_3)$ —updated belief of the fourth agent (q_4^{ABC})—for each possible combination of Alexis's, Blake's, and Chuck's decisions $[\hat{h}_1 \hat{h}_2 \hat{h}_3]$ when $c_{10} = c_{01} = 1$ and private signals are distorted by additive Gaussian noise with zero mean and unit variance. The updated belief is mostly dependent on Chuck's decision; the top four curves are for $\hat{h}_3 = 0$ and the bottom four curves are for $\hat{h}_3 = 1$.

D. Norah, the N th Agent

Norah, the N th agent, observes Y_N and $\hat{H}_1, \dots, \hat{H}_{N-1}$. Paralleling the arguments in the preceding sections, her prior belief update is a function of q_N as well as $\hat{H}_1, \dots, \hat{H}_{N-1}$ but not of q_1, \dots, q_{N-1} . Thus, we can define a general prior belief update function:

$$q_N^{AB \dots M} = U_N(q_N, \hat{h}_1, \hat{h}_2, \dots, \hat{h}_{N-1}).$$

The function U_n has a recurrence relation:

- For $n = 1$, $U_1(q) = q$.
- For $n > 1$,

$$U_n(q, \hat{h}_1, \dots, \hat{h}_{n-2}, 0) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{-\infty}^{\lambda(\tilde{q})} f_{Y_{n-1}|H}(y|1) dy}{\int_{-\infty}^{\lambda(\tilde{q})} f_{Y_{n-1}|H}(y|0) dy}},$$

$$U_n(q, \hat{h}_1, \dots, \hat{h}_{n-2}, 1) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{\lambda(\tilde{q})}^{\infty} f_{Y_{n-1}|H}(y|1) dy}{\int_{\lambda(\tilde{q})}^{\infty} f_{Y_{n-1}|H}(y|0) dy}},$$

where $\tilde{q} = U_{n-1}(q, \hat{h}_1, \dots, \hat{h}_{n-2})$.

Fig. 2 depicts the function $U_4(q_4, \hat{h}_1, \hat{h}_2, \hat{h}_3)$ for $N = 4$ for eight possible combinations of Alexis's, Blake's, and Chuck's decisions $[\hat{h}_1 \hat{h}_2 \hat{h}_3]$. An interesting property of U_n is that the updated belief is much more dependent on the most recent decision \hat{h}_{n-1} than on the earlier decisions $\hat{h}_1, \dots, \hat{h}_{n-2}$. This is especially the case when the $(n-1)$ st agent has not followed precedent. This is because the n th agent rationally concludes that the $(n-1)$ st agent observed strong evidence to justify a deviation from precedent. For example, if the decision history of the first five agents is $[0 \ 0 \ 0 \ 0 \ 1]$ then the sixth agent takes the last decision 1 seriously even though four agents chose 0. A reversal of an arbitrarily long precedent sequence may occur because we assume unbounded private signals; if private signals are bounded like in [2], [3], then the influence of precedent can reach a point where agents cannot receive a signal strong enough to justify a decision running counter to precedent.

V. OPTIMAL INITIAL BELIEF

We have constructed the prior belief update and decision making model in Section IV. In this section, we want to investigate when the system can achieve the minimum Bayes risk. For simplicity, we only consider $N = 2$. Note that the Bayes risk of the system is the same as Blake's Bayes risk because his decision is adopted as the final decision.

Let us recapitulate the computation of Blake's Bayes risk. Alexis chooses her decision threshold as $\lambda_1 = \lambda(q_1)$. Her probabilities of errors are given by

$$P_{e,1}^I = p_{\hat{H}_1|H}(1|0) = \int_{\lambda_1}^{\infty} f_{Y_1|H}(y|0) dy,$$

$$P_{e,1}^{II} = p_{\hat{H}_1|H}(0|1) = \int_{-\infty}^{\lambda_1} f_{Y_1|H}(y|1) dy.$$

Blake thinks that Alexis uses the decision threshold $\lambda_{1B} = \lambda(q_2)$ and computes her probabilities of errors differently:

$$P_{e,1B}^I = p_{\hat{H}_1|H}(1|0)_B = \int_{\lambda_{1B}}^{\infty} f_{Y_1|H}(y|0) dy, \quad (17a)$$

$$P_{e,1B}^{II} = p_{\hat{H}_1|H}(0|1)_B = \int_{-\infty}^{\lambda_{1B}} f_{Y_1|H}(y|1) dy. \quad (17b)$$

When Alexis decides $\hat{H}_1 = 0$, Blake updates his belief q_2 to q_2^0 :

$$\begin{aligned} \frac{q_2^0}{1 - q_2^0} &= \frac{q_2}{1 - q_2} \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}} \\ \implies q_2^0 &= \frac{q_2(1 - P_{e,1B}^I)}{q_2(1 - P_{e,1B}^I) + (1 - q_2)P_{e,1B}^{II}}, \end{aligned} \quad (18)$$

and his decision threshold is $\lambda_2^0 = \lambda(q_2^0)$. His probabilities of errors are given by

$$P_{e,2}^{I_0} = p_{\hat{H}_2|\hat{H}_1,H}(1|0,0) = \int_{\lambda_2^0}^{\infty} f_{Y_2|H}(y|0) dy, \quad (19a)$$

$$P_{e,2}^{II_0} = p_{\hat{H}_2|\hat{H}_1,H}(0|0,1) = \int_{-\infty}^{\lambda_2^0} f_{Y_2|H}(y|1) dy. \quad (19b)$$

Likewise, when Alexis decides $\hat{H}_1 = 1$, Blake updates his belief q_2 to q_2^1 :

$$\begin{aligned} \frac{q_2^1}{1 - q_2^1} &= \frac{q_2}{1 - q_2} \frac{P_{e,1B}^I}{1 - P_{e,1B}^{II}} \\ \implies q_2^1 &= \frac{q_2 P_{e,1B}^I}{q_2 P_{e,1B}^I + (1 - q_2)(1 - P_{e,1B}^{II})}, \end{aligned} \quad (20)$$

and his decision threshold is $\lambda_2^1 = \lambda(q_2^1)$. His probabilities of errors are given by

$$P_{e,2}^{I_1} = p_{\hat{H}_2|\hat{H}_1,H}(1|1,0) = \int_{\lambda_2^1}^{\infty} f_{Y_2|H}(y|0) dy, \quad (21a)$$

$$P_{e,2}^{II_1} = p_{\hat{H}_2|\hat{H}_1,H}(0|1,1) = \int_{-\infty}^{\lambda_2^1} f_{Y_2|H}(y|1) dy. \quad (21b)$$

Now we can compute Blake's Bayes risk R_2 :

$$\begin{aligned}
R_2 &= c_{10}p_{\hat{H}_2,H}(1,0) + c_{01}p_{\hat{H}_2,H}(0,1) \\
&= c_{10}p_{\hat{H}_2|\hat{H}_1,H}(1|0,0)p_{\hat{H}_1|H}(0|0)p_H(0) \\
&\quad + c_{10}p_{\hat{H}_2|\hat{H}_1,H}(1|1,0)p_{\hat{H}_1|H}(1|0)p_H(0) \\
&\quad + c_{01}p_{\hat{H}_2|\hat{H}_1,H}(0|0,1)p_{\hat{H}_1|H}(0|1)p_H(1) \\
&\quad + c_{01}p_{\hat{H}_2|\hat{H}_1,H}(0|1,1)p_{\hat{H}_1|H}(1|1)p_H(1) \\
&= c_{10} [P_{e,2}^{I_0}(1 - P_{e,1}^I) + P_{e,2}^{I_1}P_{e,1}^I] p_0 \\
&\quad + c_{01} [P_{e,2}^{II_0}P_{e,1}^{II} + P_{e,2}^{II_1}(1 - P_{e,1}^{II})] (1 - p_0). \quad (22)
\end{aligned}$$

The Bayes risk R_2 in (22) is a function of q_1 and q_2 . It seems natural that R_2 is minimum at $q_1 = q_2 = p_0$ because Alexis will make the best decision she can and Blake will not misunderstand her decision. Surprisingly, however, this turns out not to be true. We will prove it by studying Alexis's optimal belief q_1^* with respect to minimizing R_2 .

Let us consider the first derivative of (22) with respect to q_1 :

$$\frac{dR_2}{dq_1} = c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0}) \frac{dP_{e,1}^I}{dq_1} + c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1}) \frac{dP_{e,1}^{II}}{dq_1}.$$

We want to find q_1 that makes this first derivative zero. Using

$$\begin{aligned}
\frac{dP_{e,1}^I}{dq_1} &= \frac{dP_{e,1}^I}{d\lambda_1} \frac{d\lambda_1}{dq_1} = -f_{Y_1|H}(\lambda_1|0) \frac{d\lambda_1}{dq_1}, \\
\frac{dP_{e,1}^{II}}{dq_1} &= \frac{dP_{e,1}^{II}}{d\lambda_1} \frac{d\lambda_1}{dq_1} = f_{Y_1|H}(\lambda_1|1) \frac{d\lambda_1}{dq_1},
\end{aligned}$$

this occurs when

$$\begin{aligned}
&c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})f_{Y_1|H}(\lambda_1|0) \\
&\quad = c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})f_{Y_1|H}(\lambda_1|1) \\
&\iff \frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}. \quad (23)
\end{aligned}$$

Note that $\lambda_1 = \lambda(q_1)$ is a solution to (5),

$$\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}q_1}{c_{01}(1-q_1)}.$$

Therefore Alexis's optimal belief q_1^* needs to satisfy

$$\frac{q_1^*}{1-q_1^*} = \frac{p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}, \quad (24)$$

where the value of $(P_{e,2}^{I_1} - P_{e,2}^{I_0})/(P_{e,2}^{II_0} - P_{e,2}^{II_1})$ does not have to be 1. In fact, in additive Gaussian noise cases, it is not equal to one except for $p_0 = c_{01}/(c_{10} + c_{01})$. Therefore, the optimal value of q_1 is not p_0 in general.

VI. EXAMPLE: GAUSSIAN LIKELIHOODS

Suppose that the n th agent receives the signal $Y_n = H + W_n$ where the additive noises W_n are iid with pdf

$$f_W(w) = \frac{1}{\sqrt{2\pi}} e^{-w^2/2}. \quad (25)$$




Fig. 3. Visualization of the Bayes risk for various q_1 and q_2 for $N = 2$, $c_{10} = c_{01} = 1$, $p = 0.3$, and additive Gaussian noise with zero mean and unit variance. Alexis's and Blake's optimal prior beliefs (▲) are different from the true prior probability (●).

Each likelihood $f_{Y_n|H}$ is thus Gaussian with mean H and variance σ^2 . For a prior belief q_n , the likelihood ratio test

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \underset{\hat{H}_n(y_n)=0}{\overset{\hat{H}_n(y_n)=1}{\gtrless}} \frac{c_{10}q_n}{c_{01}(1-q_n)} \quad (26)$$

is simplified to the following comparison:

$$\underset{\hat{H}_n(y_n)=0}{\overset{\hat{H}_n(y_n)=1}{\gtrless}} \lambda_n = \frac{1}{2} + \log \frac{c_{10}q_n}{c_{01}(1-q_n)}. \quad (27)$$

Fig. 3 clearly shows that knowing true prior probability is not optimal. We have computed the performance of the sequential decision making by two agents for $c_{10} = c_{01} = 1$, $p = 0.3$, and additive Gaussian noise with zero mean and unit variance. The Bayes risk is minimum when Alexis perceives the prior probability as 0.38 and Blake perceives it as 0.23, which is marked with a triangle. For convenience, the true probability is indicated by a circle.

Figs. 4 and 5 depict the prior beliefs that the agents should have for optimal decision making. They show several common characteristics for the additive Gaussian noise model: First, the non-terminal agents (i.e., Alexis for $N = 2$ and Alexis and Blake for $N = 3$) should have belief larger than p_0 when p_0 is small and belief smaller than p_0 when p_0 is large. We call this *open-mindedness* because it is to assign higher prior belief to outcomes that are very unlikely. Second, the last agent (i.e., Blake for $N = 2$ and Chuck for $N = 3$) should have belief smaller than p_0 when p_0 is small and belief larger than p_0 when p_0 is large. This is necessary to compensate for the biases of precedent agents. Last, there is a unique point, except for $p_0 = 0$ or $p_0 = 1$, where all agents' optimal prior beliefs are the same as the true prior probability. It occurs at $p_0 = c_{01}/(c_{10} + c_{01})$. We prove this for $N = 2$ in the appendix. We also show there for $N = 2$ that the first agent should be open minded.




Fig. 4. The trend of the optimal prior beliefs for varying p_0 for $N = 2$ (Alexis and Blake). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1$, $c_{01} = 3$.

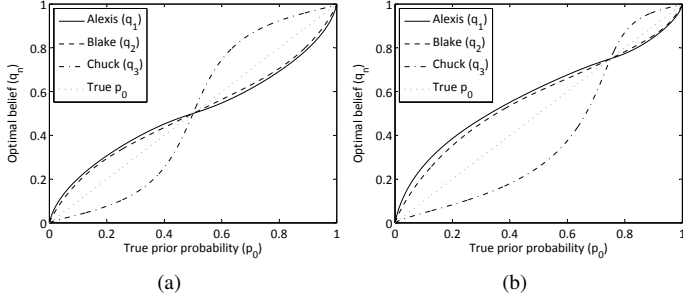


Fig. 5. The trend of the optimal prior beliefs for varying p_0 for $N = 3$ (Alexis, Blake, and Chuck). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1$, $c_{01} = 3$.

VII. CONCLUSION

We have discussed decision making sequentially performed by a group of agents that make decisions based on individually biased prior beliefs. Instead of investigating herding on a wrong action, we have assumed unbounded private signals and focused on the agents' belief update. The Bayes-optimal updated belief turns out to be the probability of each hypothesis conditioned on the decisions made by previous agents. It gets more difficult for agents in later positions to make decisions that differ from precedents; however, if one observes a very strong signal against the precedents and chooses the alternative, the decision will be taken very seriously by the following agents.

The wrong beliefs held by previous agents change the probability that following agents choose each hypothesis. Contrary to intuition, however, wrong beliefs are not always bad. In fact, the optimal beliefs of agents (those that lead to the minimum Bayes risk for the last agent) are not usually equal to the true prior probability. Specifically, in the case of observations corrupted by iid additive Gaussian noises, an agent biased toward $c_{01}/(c_{10} + c_{01})$ can be more beneficial to subsequent agents than is an accurate agent is. The point $c_{01}/(c_{10} + c_{01})$ is special because the probabilities of false alarms and missed detections will be balanced by the optimal decision rule at the prior probability. In terms of human decision making, where precedent agents are advisers or counselors to the last agent who has the final decisive power, we can say that the best advisers are necessarily open-minded people.

The idea of an open-minded adviser is related to the amount of information conveyed in the public signals. A public signal is a quantized version of a private signal while also simul-




Fig. 6. The difference between the left-hand side and the right-hand side of (29). (a) $h_1 = 1$. (b) $h_1 = 0.5$.

taneously reflecting an agent's belief. Alexis's decision will reflect her belief more than her private signal when her belief is very small (close to 0) or very large (close to 1). However, Blake would want a public signal that is most informative of Alexis's private signal. Therefore, he wants her to make her decision with a less extreme belief or an open mind.

While some conclusions of our study depend on having Gaussian likelihoods and may not hold for different types of additive noise, it is more generally true that the optimal prior beliefs are different from the true prior probability.

APPENDIX

ALEXIS'S OPTIMAL PRIOR BELIEF FOR BLAKE

For the case of $N = 2$, we investigate Alexis's prior belief that minimizes Blake's Bayes risk. Let us assume that

$$f_{Y|H}(y|0) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{y^2}{2\sigma^2}\right), \quad (28a)$$

$$f_{Y|H}(y|1) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(y - h_1)^2}{2\sigma^2}\right), \quad (28b)$$

where σ and h_1 are arbitrary positive numbers.

Conjecture 1. *If $\lambda < h_1/2$, then*

$$\begin{aligned} & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2} + \lambda h_1\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2} + \lambda h_1\right) dy \\ & < \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2} + y h_1\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2} + y h_1\right) dy. \end{aligned} \quad (29)$$

Fig. 6 depicts the difference between the left-hand side and the right-hand side of (29) for two values of h_1 and supports the conjecture. In the following we assume the conjecture to be true.

Lemma 2. *If $\lambda < h_1/2$, then*

$$\begin{aligned} & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \\ & < \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy. \end{aligned} \quad (30)$$

Proof: Substituting $y' = y/\sigma$, $\lambda' = \lambda/\sigma$, and $h'_1 = h_1/\sigma$, we obtain

$$\begin{aligned} & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \\ &= \sigma^2 \int_{-\infty}^{\lambda'} \exp\left(-\frac{y'^2}{2} + \lambda' h'_1\right) dy' \int_{\lambda'}^{\infty} \exp\left(-\frac{y'^2}{2} + \lambda' h'_1\right) dy' \end{aligned}$$

and

$$\begin{aligned} & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \\ &= \sigma^2 \int_{-\infty}^{\lambda'} \exp\left(-\frac{y'^2}{2} + y' h'_1\right) dy' \int_{\lambda'}^{\infty} \exp\left(-\frac{y'^2}{2} + y' h'_1\right) dy'. \end{aligned}$$

Since $\lambda < h_1/2$ implies $\lambda' < h'_1/2$, (30) follows from Conjecture 1. ■

Theorem 3. *Alexis's and Blake's optimal prior beliefs are equal to the true prior probability p_0 if $p_0 = c_{01}/(c_{10} + c_{01})$.*

Proof: We will show that $\partial R_2/\partial q_1 = 0$ and $\partial R_2/\partial q_2 = 0$ for $q_1 = q_2 = p_0 = c_{01}/(c_{10} + c_{01})$. Then they are Alexis's and Blake's optimal prior beliefs q_1^* and q_2^* .

First, consider $\partial R_2/\partial q_2$ using (22):

$$\begin{aligned} & \frac{\partial R_2}{\partial q_2} \\ &= -c_{10}p_0 \left[(1 - P_{e,1}^I) f_{Y_2|H}(\lambda_2^0|0) \frac{d\lambda_2^0}{dq_2} + P_{e,1}^I f_{Y_2|H}(\lambda_2^1|0) \frac{d\lambda_2^1}{dq_2} \right] \\ & \quad + c_{01}(1 - p_0) \left[P_{e,1}^{II} f_{Y_2|H}(\lambda_2^0|1) \frac{d\lambda_2^0}{dq_2} + (1 - P_{e,1}^{II}) f_{Y_2|H}(\lambda_2^1|1) \frac{d\lambda_2^1}{dq_2} \right]. \end{aligned} \quad (31)$$

From (18) and (27),

$$\lambda_2^0 = \lambda_{1B} + \log \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}},$$

and its derivative is given by

$$\begin{aligned} \frac{d\lambda_2^0}{dq_2} &= \frac{d\lambda_{1B}}{dq_2} - \frac{dP_{e,1B}^I}{dq_2} \frac{1}{1 - P_{e,1B}^I} - \frac{dP_{e,1B}^{II}}{dq_2} \frac{1}{P_{e,1B}^{II}} \\ &= \left[1 + \frac{f_{Y_1|H}(\lambda_{1B}|0)}{1 - P_{e,1B}^I} - \frac{f_{Y_1|H}(\lambda_{1B}|1)}{P_{e,1B}^{II}} \right] \frac{d\lambda_{1B}}{dq_2}. \end{aligned} \quad (32)$$

Likewise,

$$\lambda_2^1 = \lambda_{1B} + \log \frac{P_{e,1B}^I}{1 - P_{e,1B}^{II}},$$

and its derivative is

$$\frac{d\lambda_2^1}{dq_2} = \left[1 - \frac{f_{Y_1|H}(\lambda_{1B}|0)}{P_{e,1B}^I} + \frac{f_{Y_1|H}(\lambda_{1B}|1)}{1 - P_{e,1B}^{II}} \right] \frac{d\lambda_{1B}}{dq_2}. \quad (33)$$

In addition, $q_1 = q_2$ implies that $P_{e,1B}^I = P_{e,1}^I$ and $P_{e,1B}^{II} = P_{e,1}^{II}$, and we can derive the following relations for $q_1 = q_2 = p_0$:

$$\begin{aligned} \frac{f_{Y_2|H}(\lambda_2^0|1)}{f_{Y_2|H}(\lambda_2^0|0)} &= \frac{c_{10}q_2(1 - P_{e,1B}^I)}{c_{01}(1 - q_2)P_{e,1B}^{II}} = \frac{c_{10}p_0(1 - P_{e,1}^I)}{c_{01}(1 - p_0)P_{e,1}^{II}}, \\ \frac{f_{Y_2|H}(\lambda_2^1|1)}{f_{Y_2|H}(\lambda_2^1|0)} &= \frac{c_{10}q_2P_{e,1B}^I}{c_{01}(1 - q_2)(1 - P_{e,1B}^{II})} = \frac{c_{10}p_0P_{e,1}^I}{c_{01}(1 - p_0)(1 - P_{e,1}^{II})}. \end{aligned} \quad (34)$$

By substituting (32) and (33) in (31) and using the relations (34), we obtain that $\partial R_2/\partial q_2 = 0$ at $q_1 = q_2 = p_0$.

Next, we consider $\partial R_2/\partial q_1$, which is zero at q_1 and q_2 that satisfy (24),

$$\frac{q_1}{(1 - q_1)} = \frac{p_0(P_{e,2}^{I_0} - P_{e,2}^{I_1})}{(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}.$$

The condition $q_2 = c_{01}/(c_{10} + c_{01})$ leads to $\lambda_{1B} = h_1/2$ and $P_{e,1B}^I = P_{e,1B}^{II}$. Hence, $\lambda_2^0 - \lambda_{1B} = \lambda_{1B} - \lambda_2^1$ and $\lambda_2^0 = 1 - \lambda_2^1$. Then, from (19) and (21), we obtain $P_{e,2}^{I_0} = P_{e,2}^{II_1}$ and $P_{e,2}^{I_1} = P_{e,2}^{II_0}$. Therefore, only $q_2 = p_0$ completes (24) and makes $\partial R_2/\partial q_1$ zero. ■

Theorem 4. *Let $p_0 \in (0, 1)$ denote the true prior probability and q_1^* Alexis's (i.e., the first agent's) optimal prior belief.*

- If $p_0 < c_{01}/(c_{10} + c_{01})$, then $p_0 < q_1^* < c_{01}/(c_{10} + c_{01})$.
- If $p_0 = c_{01}/(c_{10} + c_{01})$, then $q_1^* = p_0$.
- If $p_0 > c_{01}/(c_{10} + c_{01})$, then $c_{01}/(c_{10} + c_{01}) < q_1^* < p_0$.

Proof: First, the proof for the case when $p_0 = c_{01}/(c_{10} + c_{01})$ is given in Theorem 3.

Next, suppose that $p_0 < c_{01}/(c_{10} + c_{01})$. Let λ in (30) denote a decision threshold according to q_2^* . Obviously, optimal prior beliefs q_1^* and q_2^* should be strictly decreasing as p_0 decreases like in Fig. 2. Hence, Theorem 3, which states that $q_1^* = q_2^* = c_{01}/(c_{10} + c_{01})$ if $p_0 = c_{01}/(c_{10} + c_{01})$, implies that

$$q_1^* < \frac{c_{01}}{c_{10} + c_{01}} \quad \text{and} \quad q_2^* < \frac{c_{01}}{c_{10} + c_{01}} \quad (35)$$

if $p_0 < c_{01}/(c_{10} + c_{01})$. Then we get $\lambda < h_1/2$ and can use (30).

Multiplying each integrand in (30) by the constant $\frac{1}{2\pi\sigma^2} \exp(-\frac{\lambda^2}{2\sigma^2} - \frac{h_1^2}{2\sigma^2})$, we get

$$\begin{aligned} & \int_{-\infty}^{\lambda} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(\lambda - h_1)^2}{2\sigma^2}\right) \exp\left(-\frac{y^2}{2\sigma^2}\right) dy \\ & \times \int_{\lambda}^{\infty} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(\lambda - h_1)^2}{2\sigma^2}\right) \exp\left(-\frac{y^2}{2\sigma^2}\right) dy \\ & < \int_{-\infty}^{\lambda} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\lambda^2}{2\sigma^2}\right) \exp\left(-\frac{(y - h_1)^2}{2\sigma^2}\right) dy \\ & \times \int_{\lambda}^{\infty} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\lambda^2}{2\sigma^2}\right) \exp\left(-\frac{(y - h_1)^2}{2\sigma^2}\right) dy. \end{aligned} \quad (36)$$

According to (28), the exponential functions in (36) are likelihood functions of Y_1 , so we have

$$\begin{aligned} & f_{Y_1|H}^2(\lambda|1) \int_{-\infty}^{\lambda} f_{Y_1|H}(y|0) dy \int_{\lambda}^{\infty} f_{Y_1|H}(y|0) dy \\ & < f_{Y_1|H}^2(\lambda|0) \int_{-\infty}^{\lambda} f_{Y_1|H}(y|1) dy \int_{\lambda}^{\infty} f_{Y_1|H}(y|1) dy. \end{aligned} \quad (37)$$

Since $\lambda = \lambda(q_2^*)$, using (17), we obtain

$$\frac{P_{e,1B}^I(1 - P_{e,1B}^I)}{f_{Y_1|H}^2(\lambda|0)} < \frac{P_{e,1B}^{II}(1 - P_{e,1B}^{II})}{f_{Y_1|H}^2(\lambda|1)}, \quad (38)$$

and (5) transforms (38) to

$$c_{10}^2 q_2^{*2} P_{e,1B}^I (1 - P_{e,1B}^I) < c_{01}^2 (1 - q_2^*)^2 P_{e,1B}^{II} (1 - P_{e,1B}^{II}).$$




Fig. 7. The point $B^0 (P_{e,2}^{I_0}, P_{e,2}^{II_0})$ always exists between the points $\bar{B}^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$ and $B^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$.

Rearranging terms gives us

$$\left(\frac{c_{10}q_2^*}{c_{01}(1-q_2^*)} \frac{1-P_{e,1B}^I}{P_{e,1B}^{II}} \right)^{-1} > \frac{c_{10}q_2^*}{c_{01}(1-q_2^*)} \frac{P_{e,1B}^I}{1-P_{e,1B}^{II}}. \quad (39)$$

The terms in the left-hand and the right-hand sides are the same as the belief update formulae (18) and (20), so this simplifies to

$$\left(\frac{c_{10}q_2^0}{c_{01}(1-q_2^0)} \right)^{-1} > \frac{c_{10}q_2^1}{c_{01}(1-q_2^1)}. \quad (40)$$

Let us discuss the meaning of the inequality (40). In Fig. 7, the convex curve depicts a flipped version of the receiver operating characteristic (ROC). When the prior belief is q , the error probabilities $(P_{e,2}^I, P_{e,2}^{II})$ are determined as the point of tangency, where the curve meets a line with slope $-c_{10}q/c_{01}(1-q)$. Two solid lines in Fig. 7 depict the lines for Blake's updated beliefs after observing $\hat{H}_1 = 0$ and $\hat{H}_1 = 1$, respectively denoted by q_2^0 and q_2^1 .

The inequality (40) restricts the range of error probabilities in which $(P_{e,2}^{I_0}, P_{e,2}^{II_0})$ can exist on the basis of $(P_{e,2}^{I_1}, P_{e,2}^{II_1})$; the point $B^0 (P_{e,2}^{I_0}, P_{e,2}^{II_0})$, a black dot, always exists on the right side of the point $\bar{B}^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$, a gray diamond. Furthermore, the point B^0 cannot exist on the right side of the point $B^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$, a black diamond, because obviously $q_2^0 > q_2^1$. Therefore, the point B^0 always exists on the curve between the points \bar{B}^1 and B^1 .

Now we draw a black dotted line that connects the points B^0 and B^1 and a gray dashed line that connects the points \bar{B}^1 and B^1 . From the restriction for the point B^0 , the slope of the former is always greater than that of the latter:

$$\frac{P_{e,2}^{II_1} - P_{e,2}^{II_0}}{P_{e,2}^{I_1} - P_{e,2}^{I_0}} > -1. \quad (41)$$

We have obtained the optimality condition (24) for Alexis's

prior belief in Section V. We can rewrite it as

$$q_1^* = \frac{p_0}{p_0 + (1-p_0) \frac{P_{e,2}^{II_0} - P_{e,2}^{II_1}}{P_{e,2}^{I_1} - P_{e,2}^{I_0}}}. \quad (42)$$

Finally, we can conclude that $q_1^* > p_0$ because of (41).

In addition, Alexis's optimal belief q_1^* is upper-bounded by $c_{01}/(c_{10} + c_{01})$ because q_1^* is strictly decreasing in p_0 and $q_1^* = c_{01}/(c_{10} + c_{01})$ when $p_0 = c_{01}/(c_{10} + c_{01})$ by Theorem 3. Combining these two bounds, we have the inequality

$$p_0 < q_1^* < \frac{c_{01}}{c_{10} + c_{01}}, \quad (43)$$

as desired.

The statement that $c_{01}/(c_{10} + c_{01}) < q_1^* < p_0$ if $p_0 > c_{01}/(c_{10} + c_{01})$ can be proven similarly. ■

ACKNOWLEDGMENT

Discussions with V. Krishnamurthy, J. Z. Sun, V. Montazerhodjat, and L. R. Varshney are greatly appreciated.

REFERENCES

- [1] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *Journal of Political Economy*, vol. 101, no. 4, pp. 612–643, Aug. 1993.
- [2] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, Aug. 1992.
- [3] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.
- [4] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.
- [5] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, Mar. 2011.
- [6] V. Krishnamurthy, "Quickest detection POMDPs with social learning: Interaction of local and global decision makers," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.
- [7] M. E. Hellman and T. M. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, vol. 41, no. 3, pp. 765–782, Jun. 1970.
- [8] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.
- [9] J. Neyman and E. S. Pearson, "On the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London. Series A*, vol. 231, pp. 289–337, Jan. 1933.
- [10] M. J. Brown, E. Henriquez, and J. Groscup, "The effects of eyeglasses and race on juror decisions involving a violent crime," *American Journal of Forensic Psychology*, vol. 26, no. 2, pp. 25–43, Sep. 2008.
- [11] J. E. Stewart, "Defendant's attractiveness as a factor in the outcome of criminal trials: An observational study," *Journal of Applied Social Psychology*, vol. 10, no. 4, pp. 348–361, Aug. 1980.
- [12] C. Stephan and J. C. Tully, "The influence of physical attractiveness of a plaintiff on the decisions of simulated jurors," *Journal of Social Psychology*, vol. 101, no. 1, pp. 149–150, Feb. 1977.
- [13] M. G. Efran, "The effect of physical appearance on the judgment of guilt, interpersonal attraction, and severity of recommended punishment in a simulated jury task," *Journal of Research in Personality*, vol. 8, no. 1, pp. 45–54, Jun. 1974.
- [14] W. K. Viscusi, "Are individuals Bayesian decision makers?" *American Economics Review*, vol. 75, no. 2, pp. 381–385, May 1985.
- [15] J. A. Swets, W. P. Tanner, Jr., and T. G. Birdsall, "Decision processes in perception," *Psychological Review*, vol. 68, no. 5, pp. 301–340, Sep. 1961.
- [16] M. Glanzner, A. Hilford, and L. T. Maloney, "Likelihood ratio decisions in memory: Three implied regularities," *Psychonomic Bulletin & Review*, vol. 16, no. 3, pp. 431–455, Jun. 2009.
- [17] G. L. Brase, L. Cosmides, and J. Tooby, "Individuation, counting, and statistical inference: The role of frequency and whole-object representations in judgment under uncertainty," *Journal of Experimental Psychology: General*, vol. 127, no. 1, pp. 3–21, Mar. 1998.

- [18] K. R. Varshney and L. R. Varshney, "Quantization of prior probabilities for hypothesis testing," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4553–4562, Oct. 2008.
- [19] J. B. Rhim, L. R. Varshney, and V. K. Goyal, "Conflict in distributed hypothesis testing with quantized prior probabilities," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2011, pp. 313–322.
- [20] —, "Benefits of collaboration and diversity in teams of categorically-thinking decision makers," in *Proc. IEEE Sensor Array and Multichannel Signal Processing Workshop*, Hoboken, NJ, Jun. 2012.
- [21] —, "Quantization of prior probabilities for collaborative distributed hypothesis testing," *IEEE Transactions on Signal Processing*, vol. 60, no. 9, pp. 4537–4550, Sep. 2012.
- [22] —, "Distributed decision making by categorically-thinking agents," in *Decision Making with Imperfect Decision Makers*, ser. Intelligent Systems Reference Library, M. Kárný, T. V. Guy, D. Wolpert, and A. E. P. Villa, Eds. Springer, 2013, to appear.
- [23] R. Radner, "Team decision problems," *Ann. Math. Stat.*, vol. 33, no. 3, pp. 857–881, Sep. 1962.
- [24] J. B. Rhim and V. K. Goyal, "Keep ballots secret: On the futility of social learning in decision making by voting," arXiv:1212.5855v1 [cs.IT], Dec. 2012.